



# HPC Storage as a Blank Canvas in Google Cloud


Dean Hildebrand

Technical Director, Office of the CTO, Google Cloud

Google Cloud



# HPC in Google Cloud



# Blank Canvas (Sheet of Ice)



# Embrace HPC Application Differences

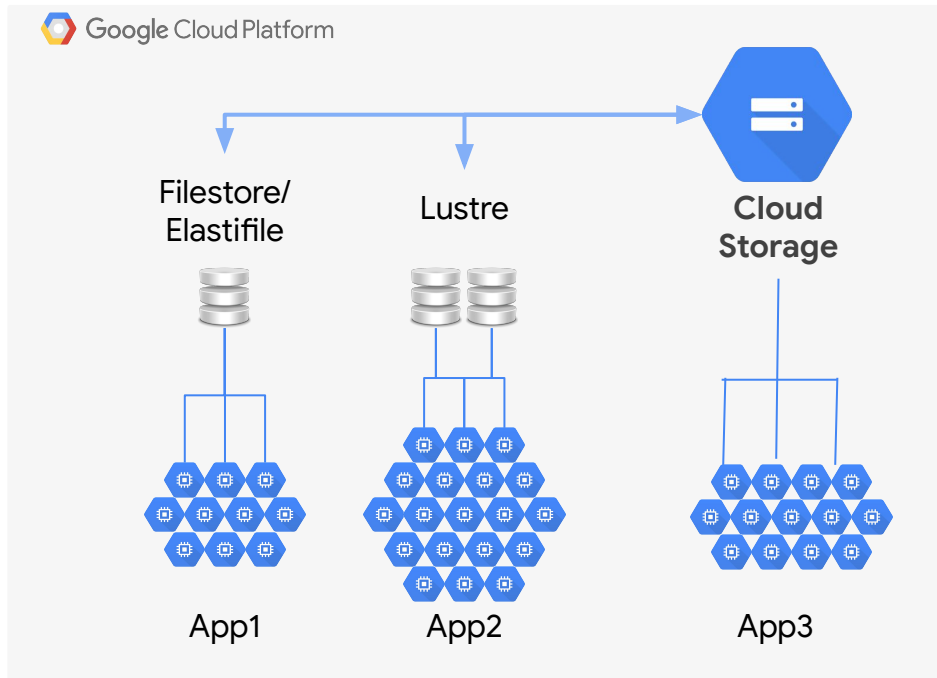
Many applications running on-prem

- Do not share datasets
- Have different I/O requirements

Large storage systems have issues

- Few applications can fully utilize them
- I/O contention is the norm
- Burst buffers difficult to integrate

Customize your file system to your application by spinning up numerous file systems with different configurations



# Google Compute Engine

## Custom Machine Types

- Size for your workload
- Recommendation engine for optimum machine size

## General Compute

- Up to 416 vCPUs
- Up to 12TB RAM
- Up to 3TB Local SSD
- Up to 8 GPUs
- Up to 7TB Intel Optane

## Compute Optimized VMs

- Fast Clocks - 3.8 GHz
- Up to 60 vCPU, 240GB RAM
- vNUMA exposure
- Overhead reduction

## Networking

- Globally routed subnets default
- **2Gbps per vCPU, up to 32Gbps**
- 15,000 VMs per Virtual Private Cloud (VPC)

## Accelerators

- Up to 8 NVIDIA GPUs
- NVIDIA K80, P100, P4, V100, T4
- NVLink and PCIe x16 for maximum performance
- Google Cloud TPUs for ML

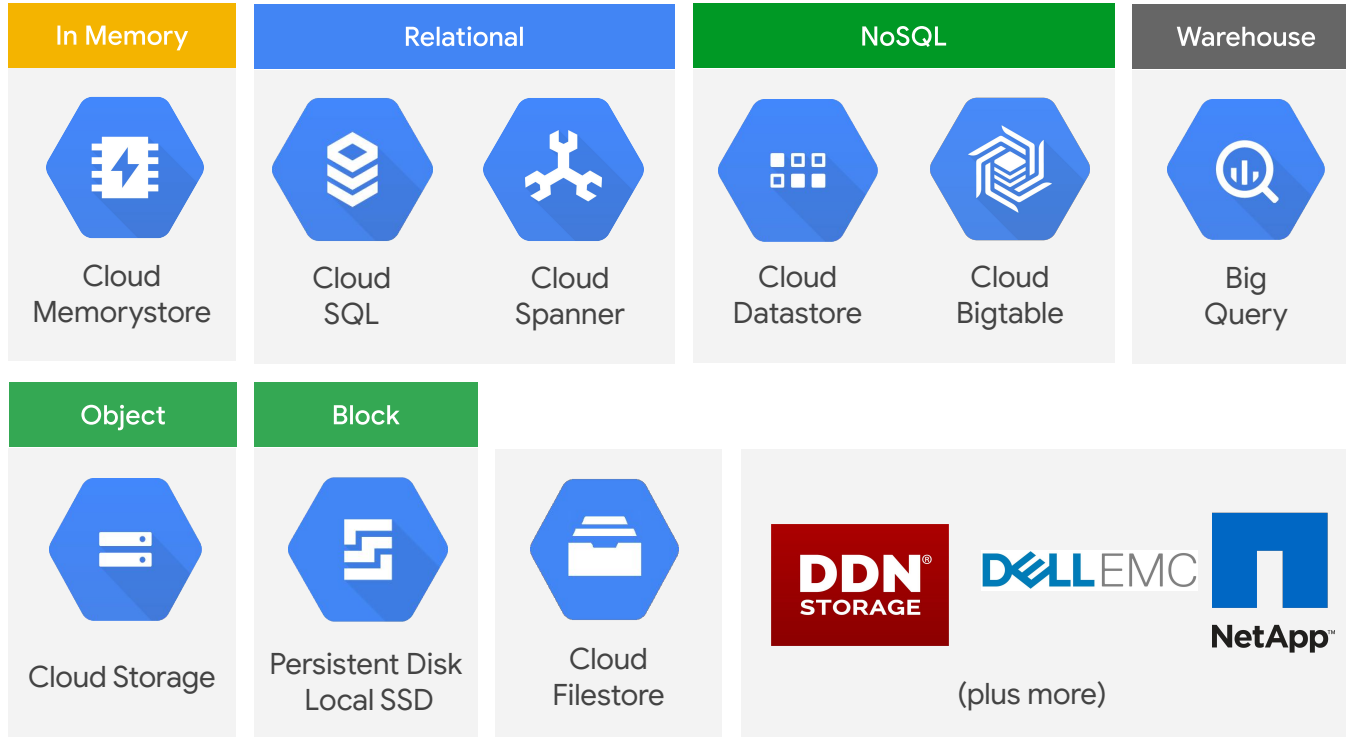
## Preemptible VMs

- Up to 80% discount
- Custom Machine Types, GPUs, Local SSDs




Google Cloud

# Storage Flexibility



# GCS Regional and Dual-Regional: Best for High Performance

- Ideal for:
  - HPC (especially when combined with Local SSD)
  - Data Analytics
  - AI/ML
- Common use cases:
  - Genomics analysis, data transcoding
  - Analytics for ecommerce and IoT
  - Integrated with Cloud DataProc, Cloud Machine Learning, BigQuery
- Best Practices
  - Remember its a general storage service
  - Ask about how caching works
  - It's all about parallelism
  - Remember you are trading latency for scale



# DDN<sup>®</sup> STORAGE Lustre in Google Cloud

## IO500 – The Only Cloud in List of Top HPC Storage Systems

### 3 spots on the IO500

- Each config demonstrates different price/perf


### Design

- Up to 200 Lustre clients
- Up to 120 Lustre storage servers
- 180TB - 1PB capacity (SSD and Disk)
- Durable and scratch configurations

### Performance

- Achieved up to 107GB/s
- Scale linearly with storage nodes and capacity

**Launch numerous file systems of any size in minutes!**





# Google Cloud Persistent Disk

Forget about devices and focus on provisioned performance

- It's all about provisioning (You get XX IOPs...period)
  - Underlying storage system has >>>> IOPs than VM allows
  - One or N disks gives same performance - no striping needed across disks
- Tiers exposed to users (but it is all abstract)
  - Disk - focus on bandwidth
  - SSD - focus on latency and IOPs
- All QoS is per VM, so it's all about what is advertised



# Defining What to Run for io500

## Current Leader

- 510GB/s Read BW
- 330GB/s Write BW
- 5763.56 Read kIOPs

PD SSD		
	per/GB	Max
Read MB/s	0.48	1200
Write MB/s	0.48	800
Read IOPS	30	100000
Write IOPS	30	30000
Price/GB/mo	\$0.17	\$108.80

Storage Type	Instance Type	Per OSS Capacity (GB)	BW \$ for month	IOPS \$ for month	\$ Max for month	Total Cap (PB)	Num OSS VMs	\$ for 1 hour
PD-Standard	n1-standard-16	10000	\$650,397	\$1,514,587	\$1,514,587	19.212	1,921	<b>\$2,104</b>
PD-SSD	n1-standard-32	3500	\$582,981	\$79,060	\$582,981	1.063	425	<b>\$810</b>
Local SSD	n1-standard-8	1500	\$74,057	\$2,663	\$74,057	0.354	236	<b>\$103</b>
Local SSD	n1-standard-8	3000	\$102,342	\$3,680	\$102,342	0.707	236	<b>\$142</b>



## Notes:

- prices will change
- Does not include discounts


# HPC file system performance

**Compute VMs**  
(File system clients)

**Virtual Private Cloud**  
(VPC)

**Storage Server VMs**

**Storage Server Storage**  
Persistent Disk  
(Standard or SSD)



- 2Gbps/vCPU network scaling
  - At least 8 vCPUs gives full 32Gbps
- Standard or Preemptible
- Customizable Machine Type
- Local SSD for Read Caching

- 2Gbps/vCPU network scaling
  - At least 16 vCPUs gives full 32Gbps
- Tiered PD-SSD vCPU requirements
  - 64 vCPUs for max performance
- Local SSD for Read Caching
  - Max performance at 1.5TB

- Performance linear to capacity
- Scales with VM, not number of PDs
- PD-Standard
  - Max bandwidth/IOPs at 10TB
- PD-SSD
  - Max bandwidth at 2.5TB
  - Max IOPs at 3.5TB



All details can be found at:  
<https://cloud.google.com/compute/docs/disks/performance>